

# OPTIMIZATION AND PARALLELIZATION OF DFT AND TDDFT IN GAMESS ON DOD HPC MACHINES

Michael E. Lasinski, Nichols A. Romero, Anthony D. Yau, Gary Kedziora, and Jean-Philippe Blaudeau\*  
High Performance Technologies, Inc.,  
Reston, VA 20190

Shawn T. Brown  
Pittsburgh Supercomputing Center, Carnegie Mellon University,  
Pittsburgh, PA 15213

## ABSTRACT

The quantum chemistry package General Atomic and Molecular Electronic Structure System (GAMESS) is employed in the first-principles modeling of complex molecular systems by using the density functional theory (DFT) as well as a number of other post-Hartree-Fock (HF) methods. Both DFT and time-dependent DFT (TDDFT) are of particular interest to the DoD Computational Biology, Chemistry, and Materials Science (CCM). Millions of CPU hours per year are expended by GAMESS calculations on DoD high performance computing (HPC) systems. Therefore, any reduction in wall-clock time for these calculations will represent a significant saving in CPU hours. As part of this work, three areas for improvement were identified: (1) replacement of the exchange-correlation (XC) integration grid, (2) TDDFT parallelization, and (3) profiling and optimization of the DFT and TDDFT codes. We summarize the work performed in these task areas and present the resulting speed-up. Many of our software enhancements are available to the general public in the 11APR2008R1 version of GAMESS with the anticipation that all of them will be available in a future release version of GAMESS.

## 1. INTRODUCTION

GAMESS (Schmidt 1993) is developed by the Mark Gordon research group at the Department of Energy Ames Lab and Iowa State University (<http://www.msg.ameslab.gov/GAMESS>). It is used for the computational modeling of complex molecular systems using a number of post-HF methods. DFT (Hohenberg 1964; Kohn 1965) and TDDFT (Casida 1995) are particularly useful because they are computationally tractable for large molecules.

As part of this work, three key task areas were identified. The first was the addition of a Lebedev (Lebedev) grid to GAMESS. This is important because

the 24MAR2007R5 version of GAMESS uses a Legendre quadrature with a high density of points at the poles compared to the equator (Murray 1993). Most high-performance DFT codes use a Lebedev sphere quadrature for the angular components of the integration which exploit the octahedral symmetry of the integrals (Lebedev 1976; Lebedev 1977; Lebedev 1992). This addition, as well as the inclusion of pruned grids such as the SG-1 grid (Gill 1993) will make the exchange-correlation (XC) calculation methods in GAMESS comparable to those in other codes. The second task area was the parallelization of the TDDFT in GAMESS. This is necessary because the 24MAR2007R5 version of GAMESS possesses a serial only implementation of TDDFT. Finally, the third task area was to optimize the DFT and TDDFT code in GAMESS so that it is more efficient. The target machines for this work are Sapphire (Cray XT3 at ERDC MSRC), Eagle (SGI Altix 3700 at AFRL MSRC) and Hawk (SGI Altix 4700 at AFRL MSRC).

## 2. DOD RESEARCH AREAS IMPACTED BY GAMESS

DFT and TDDFT calculations are of particular interest to many research efforts within the DoD. All of the aforementioned task areas will have a significant impact on these types of calculations within GAMESS. Detailed below are a few of these research areas and how they will benefit from our improvements.

The Army Institute for Multi-Scale Reactive Modeling (MSRM), awarded to the Army Research Laboratory and lead by Betsy Rice (ARL/WMRD), is interested in developing computational techniques that reduce the empiricism in modeling the sensitivity of DoD munitions. These techniques will incorporate fundamental physics and chemistry methods that bridge the gap between atomistic, meso, and continuum scales. Ultimately, the Institute's software will reduce the risk, cost, and time to develop new munitions that comply with the DoD's Insensitive Munitions directive.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE <b>DEC 2008</b>		2. REPORT TYPE <b>N/A</b>		3. DATES COVERED <b>-</b>	
4. TITLE AND SUBTITLE <b>Optimization And Parallelization Of DFT And TDDFT IN GAMESS on DOD HPC Machines</b>				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) <b>High Performance Technologies, Inc., Reston, VA 20190</b>				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT <b>Approved for public release, distribution unlimited</b>					
13. SUPPLEMENTARY NOTES <b>See also ADM002187. Proceedings of the Army Science Conference (26th) Held in Orlando, Florida on 1-4 December 2008, The original document contains color images.</b>					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT <b>UU</b>	18. NUMBER OF PAGES <b>8</b>	19a. NAME OF RESPONSIBLE PERSON
a. REPORT <b>unclassified</b>	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>			

GAMESS is one of the ab-initio wavefunction codes that will be used to parameterize the Reax Force Field used in reactive molecular dynamics at the atomistic scale. The proposed improvements to GAMESS will provide a faster DFT code that will give researchers access to larger problems that will be critical to providing high-quality reference data at the molecular level. (<http://hpcmo.hpc.mil/community/SAS/institutes.php>)

The need to reduce the amount of weight devoted to dry mass (e.g. fuel tanks) of future spacecraft has motivated the research of novel high energy density materials (HEDM) with a large specific impulse. These chemical propellants would offer greater thrust per unit density and would permit larger payloads to be launched into orbit reducing the overall space deployment costs. Additionally, the extreme temperatures encountered by spacecraft have led to the search for protective coatings with enhanced thermal and physical properties. A potential group of candidate molecules are the polyhedral oligomeric silsesquioxanes (POSS). Due to their chemical nature, they show great promise as additives into common plastics via co-polymerization and blending. As very little is known about the synthesis of POSS compounds, computational chemistry methods, such as DFT, can greatly contribute to the identification of candidate molecules. As HEDM and POSS compounds are large molecules with high Z elements, expensive DFT calculations are typically required. Jerry Boatz (AFRL/PRSP) is a DoD scientist with a Challenge Project allocation to investigate these two classes of materials. The implementation of the Lebedev grid, pruned grids, and DFT optimizations will drastically reduce the time-to-solution for these calculations. (<http://www.arsc.edu/challenges/materials.html>)

Nonlinear optical (NLO) materials are of great interest due to their capability of converting laser light (typically mid- and far-infrared) to light with longer or shorter wavelengths. Porphyrin-based molecules show promise as optical limiting media for use in infrared counter measures on aircraft. Ruth Pachter (AFRL/MLPJ) is the lead DoD researcher applying TDDFT to calculations of the absorption spectra of candidate molecules. Understanding the linear absorption spectra of chromophores is a necessary step in understanding nonlinear phenomena of many materials. The parallelization of the TDDFT in GAMESS would allow the calculation of absorption spectra of larger molecules and expedite time-to-delivery for DoD researchers investigating NLO materials for use in a number of applications spanning the gamut of organic semiconductors to optical storage devices. (<http://www.arsc.edu/challenges/materials.html>)

The ever growing need for energy has motivated research into alternative energy sources and into efficient methods of energy storage. Of the many different scenarios for energy production and storage being examined, there are certain areas of commonality amongst almost all of them. Whether considering solar cells, fuel cells, capacitors, or batteries, most of the interesting processes occur at the interfaces between dissimilar materials, often involving transfer of charge. Such electrochemical processes have been used by humans in one form or another for hundreds, possibly thousands of years. However, our understanding of these electrochemical processes at the atomic and mesoscale is still fairly rudimentary. Douglas Dudis and Todd Yeates (AFRL/ML) are performing computational and experimental research to help elucidate some of the processes that are manifested in electrochemical devices. (Yeates, personal communication, 2007) This often involves a density functional description of materials in interfacial geometries – an inherently nanoscale problem. Faster DFT and TDDFT routines in GAMESS will help provide more insight into these electrochemical processes more quickly than would otherwise be possible.

### 3. DENSITY FUNCTIONAL THEORY

DFT (Hohenberg 1964) is the foundation for modern day first-principles calculations on atoms, molecules, surfaces, and condensed matter systems. It is a set of existence theorems relating the ground state many-body  $\psi(\{r_i\})$  wave function to its charge density  $n(r)$ . DFT does not prescribe how to obtain  $\psi(\{r_i\})$  nor the physical quantities associated with it for an arbitrary system. It is the work of Kohn and Sham (KS) (Kohn 1965) that provided the theoretical framework for actual calculations. The essence of the KS approach is to replace the  $\psi(\{r_i\})$  by an auxiliary system with an effective Hamiltonian  $H_{KS}$  and single-particle wave functions  $\{\psi_i(r)\}$ . The many-body effects are instead treated in a mean-field manner with all the quantum mechanical correlation effects approximated by an XC functional  $E_{xc}[n(r)]$ . The KS-DFT ground state in GAMESS is obtained by a self-consistent field (SCF) solution of the KS equation,

$$\left[ \frac{\hbar^2}{2m} \nabla^2 + V_{ext}(r) + e^2 \int \frac{n(r')}{|r-r'|} d^3r' + V_{XC}(r) \right] \psi_i(r) = \epsilon_i \psi_i(r) \quad (1)$$

$V_{eff}(r)$

where the terms on the left-hand side of Eqn. (1) are, from left to right, the single-particle kinetic energy, the

external potential from the nuclei, the classical Coulomb energy, and  $V_{xc}(r)$  which equals  $\delta E_{xc}[n(r)]/\delta n(r)$ . The SCF procedure consists of calculating  $V_{eff}(\gamma)$  for a given initial  $\{\psi_i(r)\}$ , constructing  $H_{KS}$ , and a diagonalization step to determine the next set of  $\{\psi_{i+1}(r)\}$  and  $\{\varepsilon_{i+1}\}$ . This procedure is repeated until a SC solution is obtained within a pre-defined tolerance.

GAMESS expands  $\psi_i(r) = \sum_{\mu} c_{i\mu} \phi_{\mu}(r)$  in a Gaussian basis set; permitting the calculation of the kinetic and potential energy terms in Eqn. (1) using analytic, albeit complex, formulae. The one exception is  $V_{xc}(r)$  which is calculated numerically on atom-centered grids. The computation of the XC term scales as  $N_{grid}N_{basis}^2$ , where  $N_{grid}$  is the number of grid points and  $N_{basis}$  is the number of basis functions.

The SCF procedure in GAMESS typically consists of a few computationally cheap HF SCF steps which precondition the initial DFT guess. The DFT SCF solution procedure consists of two segments which are calculated respectively on a coarse and then on a fine grid. A brief summary of all of the grids now available in GAMESS is the subject of the next section.

#### 4. LEBEDEV AND SG-1 GRID

The computation of  $E_{xc}[n(r)]$  and its associated quantities involves 3D molecular integrals of the form shown in Equation 2.

$$\int F(n(r), \nabla n(r), \dots) d^3r \quad (2)$$

As  $F(n(r), \nabla n(r), \dots)$  is generally quite complicated it is not possible to calculate these integrals analytically. Becke showed that these 3D integrals can be accurately approximated by sums of one-center atomic-like integrals centered on each nuclei (Becke 1998). These atom-centered integrals can then be computed by well-known numerical quadrature schemes. GAMESS uses an Euler-Maclaurin (EM) quadrature for the radial grid and a Gauss-Legendre (GL) direct product quadrature for the  $(\theta, \phi)$ -grid (Murray 1993). The quadrature parameters in GAMESS are specified using the keywords NRAD0, NTHETA0, & NPHI0 for the initial coarse grid and NRAD, NTHETA, & NPHI for the fine grid. The Lebedev sphere quadrature uses approximately 1.5 times fewer points than a GL quadrature of comparable accuracy. The Lebedev points and their associated weights are generated using the code of Chipman (2002) which is based on the code of Laikov (Lebedev 1999). Similar to the existing grid, Lebedev angular grids, up to a value of

5810, can be specified using the keywords NLEB0 and NLEB. We denote the EM-GL and EM-Lebedev grid parameters by specifying two and three parameters, respectively, in parentheses.

The pruned grid of Gill *et al.* (Gill 1993), called SG-1, was also implemented as part of this work. Pruned grids take advantage of the variability of the electron density with respect to the distance from the nuclei by specifying the density of angular quadrature as a function of location on the radial quadrature. The total quadrature grid is the sum of the atom-centered grids at each radial quadrature point. In direct product grids, angular quadratures with the same number of points are used at each radial point. This leads to a large number of unnecessary quadrature points, particularly very close to and far away from the nucleus. The SG-1 grid is constructed by five different levels of angular quadrature, the smallest being a Lebedev grid of 6 points and the largest having 194 points (Gill 1993) with accuracy comparable to the (50,194) EM-Lebedev grid. The resultant grid has approximately 3,000 total points per atom versus the 9,700 points for the EM-Lebedev grid or 14,400 points for the EM-GL grid.

Two additional pruned grids were implemented as part of this work. These grids were denser angular grids than the SG-1. As a result of having more points, they are also more accurate. They were provided to us by Curtis Janssen (2008, personal communication). The first grid has 95 radial shells and varies from 6 to 434 angular points. The second pruned grid has 155 radial shells and varies from 86 to 974 angular points. In both cases, the variation in angular points is similar to that of the SG-1 grid as described by Gill *et al.* (Gill 1993).

#### 5. PROFILING AND OPTIMIZATION

Our profiling efforts focused on a DFT solvation calculation for the TPABr<sub>3</sub> molecule, an organic dye used in lasers. This input file was provided by Douglas Dudis and Todd Yeates of AFRL/RX. Its modest size allowed it to easily fit on many of the debug queues on the target machines. This input file was profiled using the Tuning and Analysis Utilities (TAU) Performance System<sup>TM</sup> (University of Oregon, <http://www.cs.uoregon.edu/research/tau/home.php>) and the PAPI tools (Innovative Computing Laboratory at the University of Tennessee at Knoxville, <http://icl.cs.utk.edu/papi>). TAU is a profiling and tracing toolkit for performance analysis of parallel programs and the PAPI library is used to gather events from the hardware counters, e.g. cache misses and branch mispredictions. We observed a large number of level two

cache misses on Sapphire for this GAMESS calculation using the TAU Performance System and PAPI tools. Excessive cache misses were identified in three subroutines: DFTTTRF, DFTTTRFG, & DFTGDV. These cache misses were caused by out-of-stride memory accesses to several matrices. Other user input files which were profiled identified the same computational bottlenecks confirming these results.

The subroutines DFTTTRF and DFTTTRFG compute  $n(r)$  and  $\bar{V}n(r)$ , respectively. Cache misses were reduced by an order of magnitude by simply transposing the matrix storing the molecular orbital expansion coefficients  $C_{i\mu}$ . The third subroutine DFTGDV is used to compute nuclear gradients. Cache misses were reduced by two orders of magnitude by simply swapping indices in the code that accesses the symmetric density matrix. The speed-ups resulting from these reductions in cache misses are system-size dependent, with larger systems exhibiting greater speed-ups. For a typical DFT nuclear gradient calculation, the time-to-solution has been reduced by about a factor of 1.5 – 2.0 due to this memory-stride optimization.

## 6. TDDFT PARALLELIZATION

The formal foundations of time-dependent density functional (TDDFT) were provided by Runge and Gross in a similar manner to its time-independent analogue (Runge 1984). As with DFT, KS methodology is applied to determine the response of the electron density to a time-dependent perturbation (Casida 1995). The many-body time-dependent Schrödinger equation is replaced by a single-particle time-dependent KS equation whose solution requires the computation of two-electron integrals and XC terms the details of which can be found elsewhere (Casida 1995) and are omitted here for the sake of brevity.

The TDDFT capabilities in GAMESS were implemented by Chiba *et al.* (Chiba 2006a) and closely follow the algorithm presented in Stratmann *et al.* (Stratmann 1998). It is capable of calculating excited state energies, transition frequencies, and oscillation strengths, which are used in characterizing NLO materials. In the 24MAR2007R5 version of GAMESS, TDDFT energy and gradient calculations could only be performed serially. As a result of TAU profiling, it was determined that the TDDFT energy calculation spent most of its time performing the two-electron integrals and obtaining the XC contribution. It was also discovered that many GAMESS DFT parallel summation calls are all-reductions. Therefore, after most parallel summations in the DFT, the final value would be sent out

to each processor. Due to the similarities between the DFT and the TDDFT code, it was possible to leverage the already existing DFT code to parallelize the TDDFT. For example, the TDDFT two-electron integral subroutines were very similar to the DFT two-electron integral subroutines. Therefore, it was possible to adapt the DFT two-electron integral parallelization approach to the analogous TDDFT subroutine (TDTWOEI). For the XC contribution, the grid points were already broken up and serially calculated in groups of 200. As before, the DFT approach was modified to suit the TDDFT. In this case, there were several examples in the DFT of parallel code that looped over an array of items in a parallel manner. This approach was adapted for the XC contribution in the subroutine TDFXCP to have each processor work on different groups of 200 grid points.

The TDDFT gradient calculation, which follows the approach of Furche and Ahlrichs (Furche 2002), and implemented by Chiba *et al.* (Chiba 2006b), had two significant bottlenecks. The first came from a subroutine that eventually calls both the TDDFT two-electron integral subroutines and the XC subroutines (TDTWOEI & TDFXCP). These were effectively already in parallel due to the work done for the TDDFT energy calculation. The other bottleneck was from a subroutine, VFEXC, that was structured similarly to the XC contribution calculations. Therefore, the same approach for looping over the array of grid points was used to parallelize this subroutine.

## 7. DFT RESULTS

The impact of the Lebedev grid and the optimizations to the DFT code were quantified by using a modified version of an input file derived from the DoD High Performance Computing Modernization Program’s (HPCMP) Technology Insertion benchmark suite. This is referred to as the TI-xx standard input. (The TI-06, TI-07, and TI-08 standard input file are identical and are simply referred to as TI-xx.) In the input file, the GAMESS keyword, GUESS=HUCKEL instead of GUESS=MOREAD was used to ensure that the number of SCF iterations would be similar for all grids. Four input files were derived from the TI-xx standard input file: Default, Army, Default MINI, and Army MINI. They were constructed using two different sets of grid parameters: Default, “Army grade”; and two different sets of basis functions: MINI, 6-311G(d,p)). The intent of these four input files is to assess the performance of the new 11APR2008R1 version of GAMESS relative to the old 24MAR2007R5 version for different problem sizes.

These four input files change the size of the problem by varying both the number of grid points and the number of basis functions. In addition to these four input

files, an input file using the SG-1 pruned grid and the 6-311G(d,p) basis set was also created.

Table 1: Grid parameters for default and Army grade cases.

	EM-GL (radial, theta, phi)		EM-Lebedev (radial, angular)	
	Coarse	Fine	Coarse	Fine
Default	(24,8,16)	(96,12,24)	(24,86)	(96,194)
Army grade	(48,12,24)	(96,36,72)	(48,194)	(96,1202)

The first input file, which is referred to as Default, uses the default EM-GL grid parameters from the 24MAR2007R5 version of GAMESS. The 6-311G(d,p) Pople basis set is used here. Default is the input file most similar to the original TI-xx standard input benchmark. We also constructed a default EM-Lebedev grid with roughly 1.5 fewer points than the default EM-GL grid. This was *not* the default Lebedev grid parameters in the 11APR2008R1 version of GAMESS, but was instead chosen due its comparable accuracy to the default EM-GL grid parameters. Secondly, the Army grade (referred to simply as Army) input file was constructed in a similar fashion for both the EM-GL and EM-Lebedev grid using a high density of grid points with an integration error of about one microHartree/atom. All of the grid parameters are shown in Table 1.

The third and fourth input files use a significantly smaller basis set obtained by setting the keyword GBASIS=MINI. While the first two input files use a 6-311G(d,p) basis set, these input files use a minimal basis set (Huzinaga 1984). The third input file uses the default grids and the minimal basis set, while the fourth input file uses the “Army grade” grid and the minimal basis set. The expectation is that the third input file is representative of a very small problem, and that the fourth is representative of an intermediate sized problem.

DFT calculations for all four of these input files were performed on Eagle, Hawk, and Sapphire for cases ranging from 1 to 128 processing elements (PEs). On each machine, three different cases were performed for each input file: (1) 11APR2008R1 version of GAMESS with an EM-Lebedev grid, (2) 11APR2008R1 version of GAMESS with an EM-GL grid, and (3) 24MAR2007R5 version of GAMESS with an EM-GL grid. (Certain cases could not be obtained due to queue limits). This was done to compare the speed-up of the new 11APR2008R1 code with that of the older 24MAR2007R5 code. Due to space considerations, we are only able to show selected figures exemplifying our benchmarks.

In Figure 1, the speed-up shown is defined as the wall-clock time ratio of the 24MAR2007R5 version of

GAMESS using the EM-GL grid to that of the 11APR2008R1 version using the EM-Lebedev grid. Based on Figure 1, the Army input file case (diamonds) exhibits the largest speed-ups, over a factor of four in some cases. The Army MINI demonstrated the next largest speed-up of about a factor of two to three. The smallest speed-ups were for the Default and Default MINI cases with speed-ups of around 1.2 – 1.5. The Army and Army MINI inputs are representative of the types of calculations performed by DoD users. Naturally, the benefits of the EM-Lebedev grid and the optimizations will be greater for DoD users with larger problem sizes.




Figure 1: Speed-up of wall-clock time vs. number of processing elements for the standard TI-xx derived input files (ratio of 24MAR2007R5x EM-GL grid wall-clock time to 11APR2008R1 with EM-Lebedev grid calculation Time).

## 8. PRUNED GRID RESULTS

In order to quantify the efficiency gains obtainable with the new pruned grids, DFT calculations were performed on an AZT molecule for the EM-GL grid, two different sized Lebedev grids, and all of the pruned grids. This

was done in order to compare both the accuracy (in terms of the ground state energy) and the calculation time for all of the cases. The results shown below in Table 2 are from serial calculations done on Pople, the Pittsburgh Supercomputing Center SGI Altix 4700. The most interesting column in the table is the final column which shows two numbers. The first is the speed-up of the grid for that row of the table relative to the Army grade EM-GL grid as defined by the GAMESS developers (as shown in the first row of Table 2). This is obtained by dividing the Army grade EM-GL grid calculation time by the calculation time for the grid in that row of the table. The second number is the ratio of grid points per atom in the Army grade EM-GL grid vs. the number of grid points per atom for the grid in that row of the table. Since the computational scaling of the numerical quadrature procedure should scale linearly with the number of grid points, these numbers would be equal if the energy computation only consisted of the quadrature part of the procedure. In the case of the larger grids, the timings are essentially dominated by this step, and these two numbers are almost equal. However for the smaller grids, the other parts of the calculation are more prominent, and therefore the two numbers differ significantly. Regardless, there is a significant reduction in the computational time for the pruned grids, with the Janssen1 and Janssen2 grids giving good accuracy for the cost.

## 9. TDDFT RESULTS

For the TDDFT benchmarks, energy and gradient calculations were performed on a 7-diethylamino-4-trifluoro methyl Coumarin molecule ( $C_{14}H_{14}F_3NO_2$ , or commonly referred to as Coumarin 152A) provided by the GAMESS developers. Figure 2 shows the parallel speed-up on Sapphire for the Coumarin 152A calculation. The entire calculation (dark blue line with diamond points) scales well through 64 PEs with an

efficiency of 77%. However, the efficiency drops to 62% at 128 PEs. This is due to the TDDFT gradient (blue line with circular points) which has a lower efficiency of only 41% at 128 PEs, in sharp contrast to the TDDFT energy (green line) which still has an efficiency of 89% (speed-up of 114) at 128 PEs. These results suggest that the TDDFT energy calculation scales as well as if not better than the DFT energy calculation (red line).

In addition to the above results, a larger problem (in terms of basis functions) was developed from the previous problem increasing the basis from a Dunning/Hay double zeta basis to a 6-311G(d,f,p) basis. This resulted in a significant increase in the number of Cartesian Gaussian basis functions from 370 to 1082. The number of basis set shells increased from 162 to 278. This input file was run on Sapphire over the range of 16 to 512 PEs, and on Hawk for 256 and 500 PEs yielding similar energies and gradients on both architectures. Smaller runs on Sapphire were not done due to queue constraints.

On Sapphire, total execution time drops from 41 hours at 16 PEs to 2 hours at 512 PEs. Figure 3 shows the parallel speed-up for this large input file. It should be noted that this speed-up was obtained by extrapolating the time for 16 PEs to the time for one PE by multiplying by 16. The extrapolated serial time was then used to calculate the speed-up. Both the TDDFT energy calculation and the TDDFT gradient calculation are significantly above 50% efficiency. The DFT calculation scales above 50% efficiency through 256 PEs, but then drops below it by 512 PEs. It should also be noted that the normalization to 16 PEs means that the speed-up has been calculated using an extrapolated serial time that is most likely an upper bound to the true serial time. Therefore, the actual 512 PE speed-ups are probably slightly less than what is shown in Figure 3. However, the qualitative results remain the same.

Table 2: Computed energy, number of grid points per atom, and the timings for 6-31G\*/BLYP AZT with various grid types implemented for this work. The last column is the displays the speedup for the overall computation of the energy and the ratio of grid points per atom in the army-grade grid and the used grid.

Grid Type	E (Hartree)	Points per Atom	Time (Seconds)	Speedup/GP Ratio
EM-GL (96, 36, 72)	-962.8984441515	217728	8123	1.0/1.0
EM-Lebedev (96, 1202)	-962.8984476069	115392	3930	2.1/1.8
EM-Lebedev (96,302)	-962.8983262093	28992	1198	6.8/7.5
SG-1 (50,max 194)	-962.8985137735	~3700	407	20.0/59.0
Janssen1 (95,max 434)	-962.8984656448	~15000	774	10.5/14.5
Janssen2 (144,max 974)	-962.8984489890	~71000	2376	3.4/3.1




Figure 2: Parallel speed-up of Coumarin 152A gradient calculation on Sapphire with an Army grade Lebedev grid.




Figure 3: Large problem parallel speed-up for Coumarin 152A TDDFT gradient calculation on Sapphire. (1-8 PE speed-up extrapolated from 16 processor normalization)

## CONCLUSIONS

A number of software enhancements were implemented in GAMESS and are now available in the 11APR2008R1 release. These modifications are: 1) addition of a Lebedev grid and SG-1 pruned grid 2) TDDFT parallelization and 3) optimizations to the DFT & TDDFT code. Additional enhancements including further optimizations of the DFT and TDDFT code, and two more pruned grids are anticipated to be part of future release versions of GAMESS. Benchmarks of a Coumarin 152A molecule demonstrated comparable parallel scaling between the DFT & TDDFT calculations. Four input files were derived from the TI-xx standard benchmark and were used to quantify the speed-ups resulting from the new grid and our memory-stride optimizations. Based on these benchmarks, *all* users will find a reduction in their time-to-solution for DFT and TDDFT calculations. For users performing accurate DFT calculations on large molecules there will be a speed-up of over a factor of three in their time-to-solution. As a result, these improvements will greatly impact a number of DoD research areas, particularly DFT calculations of energetic materials, DFT

calculations of novel propellant formulations, and TDDFT calculations of NLO materials.

## ACKNOWLEDGEMENTS

This project, CCM-KY6-002, was funded by the DoD High Performance Computing & Modernization Program (HPCMP) under the User Productivity Enhancement and Technology Transfer (PET) program. We acknowledge a number of useful discussions with the PET support staff and the members of Mark Gordon's group at Iowa State University. In particular, we are very thankful to Dr. Sameer Shende of ParaTools, Inc. for helping us to interface GAMESS with TAU Performance System™ and PAPI on a number of DoD HPC systems.

## REFERENCES

- Becke, A. D., 1988: A multicenter numerical integration scheme for polyatomic molecules. *J. Chem. Phys.*, **88**, 2547-2553.
- Casida, M. E., 1995: Chapter 5: Time-dependent density functional response theory for molecules. *Recent Advances in Density Functional Methods*, Part I, D. P. Chong, Ed., World Scientific, 155-192.
- Chiba, M., Tsuneda, T., and Hirao, K., 2006a: An efficient state-specific scheme of time-dependent density functional theory. *Chem. Phys. Lett.*, **420**, 391-396.
- Chiba, M., Tsuneda, T., and Hirao, H., 2006b: Excited state geometry optimizations by analytical energy gradient of long-range corrected time-dependent density functional theory. *J. Chem. Phys.*, **124**, 144106-1-144106-11.
- Chipman, D. M. and Dupuis, M., 2002: Implementation of solvent reaction fields for electronic structure. *Theoret. Chem. Acc.*, **107**, 90-102.
- Furche, F. and Ahlrichs, R., 2002: Adiabatic time-dependent density functional methods for excited state properties. *J. Chem. Phys.*, **117**, 7433-7447.
- Gill, P. M. W., Johnson B. G., and Pople, J. A., 1993: A standard grid for density-functional calculations. *Chem. Phys. Lett.*, **209**, 506-512.
- Hohenberg, P. and Kohn W., 1964: Inhomogeneous electron gas. *Phys. Rev. B*, **136**, B864-B871.
- Huzinaga, S., Andzelm, J., Klobukowski, M., Radzio-Andzelm, E., Sakai, Y., and Tatewaki, H., 1984: *Gaussian Basis Sets for Molecular Calculations*. Elsevier, 1984, 426 pp.
- Kohn, W. and Sham, L. J., 1965: Self-consistent equations including exchange and correlation effects. *Phys. Rev.*, **140**, A1133-A1138.
- Lebedev, V. I., 1975: Values of nodes and weights in Gauss-Markov types of quadrature formulae for



- spheres from 9<sup>th</sup> to 17<sup>th</sup> order of accuracy, invariant with respect to an octahedral group with inversion. *Z. Vychisl Mat. Mat. Fiz.* **15**, 48-54.
- Lebedev, V. I., 1976: Quadratures on a sphere. *Z. Vychisl Mat. Mat. Fiz.* **16**, 293-306.
- Lebedev, V. I., 1977: Quadrature formulas for the sphere of 25<sup>th</sup> to 29<sup>th</sup> order accuracy. *Sib. Mat. Zh.* **18**, 132-142.
- Lebedev, V. I., and Skorokhodov, A. L., 1992: Quadrature formulas for a sphere of orders 41, 47, and 53. *Russ. Acad. Sci. Dokl. Math.* **45**, 587-592.
- Lebedev, V. I. and Laikov, D. N., 1999: A quadrature formula for the sphere of the 131<sup>st</sup> algebraic order of accuracy. *Dokl. Math.*, **59**, 477-481.
- Murray, C. W., Handy, N. C., and Laming, G. L., 1993: Quadrature schemes for integrals of density functional theory. *Mol. Phys.*, **78**, 997-1014.
- Runge, E. and Gross, E. K. U., 1984: Density-functional theory for time-dependent systems. *Phys. Rev. Lett.*, **52**, 997-1000.
- Schmidt, M. W., Baldridge, K. K., Boatz, J. A., Elbert, S.T., Gordon, M. S., Jensen, J. H., Koseki, S., Matsunaga, N., Nguyen, K. A., Su, S. J., Windus, T. L., Dupuis, M., and Montgomery J. A., 1993: General Atomic and Molecular Electronic-Structure System. *J. Comput. Chem.*, **14**, 1347-1363.
- Stratmann, R. E., Scuseria, G. E., and Frisch, M. J., 1998: An efficient implementation of time-dependent density-functional theory for the calculation of excitation energies of large molecules. *J. Chem. Phys.*, **109**, 8218-8224.